



Agency Use of Artificial Intelligence

Ad Hoc Committee on Agency Use of Artificial Intelligence

Draft Statement for Ad Hoc Committee | November 18, 2020

1 Artificial intelligence (AI) techniques are changing how government agencies do their
2 work. Advances in AI hold out the promise of lowering the cost of completing government tasks
3 and improving the quality, consistency, and predictability of agency decisions. But enhanced
4 agency use of AI also raises important issues, such as: how agencies might best design
5 algorithms; the risk that agencies' adoption of AI will create or exacerbate undesired biases; the
6 appropriate spheres of human and AI decisionmaking in administrative processes; the
7 compatibility of AI decisionmaking with foundational administrative law concepts like
8 transparency, accountability, and reasoned decisionmaking; and the need for adequate oversight
9 of AI decisionmaking.

10 Throughout its existence, the Administrative Conference has addressed similar questions
11 about agencies' operations and processes. Consequently, many of the Administrative
12 Conference's past recommendations may help agencies in deciding how to make the best use of
13 AI.¹ But AI also presents agencies with new questions, or requires agencies to address existing
14 questions in materially distinct contexts. That is why the Administrative Conference has
15 developed this Statement. It offers a framework for agencies to use in thinking through some of
16 the important questions presented by agencies' uses of AI. The Statement draws on a pair of
17 reports commissioned by the Administrative Conference, as well as the input of AI experts from

¹ See, e.g., Admin. Conf. of the U.S., Recommendation 2018-5, *Public Availability of Adjudication Rules*, 84 Fed. Reg. 2,142 (Feb. 6, 2019); Admin. Conf. of the U.S., Recommendation 2016-5, *The Use of Ombuds in Federal Agencies*, 81 Fed. Reg. 94,316 (Dec. 23, 2016); Admin. Conf. of the U.S., Recommendation 2015-4, *Designing Federal Permitting Programs*, 80 Fed. Reg. 78,164 (Dec. 16, 2015).

Commented [A1]: Some commenters have suggested that the whole statement has become too long and detailed.

Commented [A2]: Commenters disagree about whether to define AI in this statement. They also disagree about whether to include language about the scope of the AI being addressed in this document (e.g., to clarify whether the statement covers both learning and rule-based systems).

Relatedly, one commentator raises the following concern:

"A central, overarching concern we raise is that the statement does not distinguish between standalone AI products and embedded AI products (like Siri in iPhones, for example). In this way, it assumes that the only way agencies use or interact with AI is through distinct products or algorithms, when it is in fact very common for agencies to interact with AI embedded in commercial products. Importantly, agencies are likely to have significantly less control over the design, development, and, to some extent, use of embedded AI technologies, and thus different considerations should apply.

One way to handle this issue is to include a global statement at the beginning noting the distinction and explaining that agencies' ability to exercise control over AI embedded within commercial products may be limited (but still must comply with otherwise applicable laws, regulations, and policies)."



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

18 government, academia, and the private sector whose shared experiences suggest some pervasive
19 issues in agency use of AI.

20 The Statement highlights some of the important issues confronting agencies in their use
21 of AI.² It identifies principles agencies should consider and apply in deciding when to use AI,
22 what type of AI to use, and how to develop, implement, and maintain a given AI technique.
23 Because agencies use AI in myriad and diverse ways, this Statement does not delve into specific
24 AI techniques. Instead, it focuses on broad considerations that should inform a wide variety of
25 agency uses of AI.

26 The considerations addressed in this Statement implicate law, policy, procurement,
27 finances, human resources, and technology. It is important, therefore, that agencies involve
28 officials from all relevant offices in formulating responses to the considerations addressed in this
29 Statement.

30 Further, many of the considerations addressed in this statement are interrelated. For
31 example, accountability, judicial review, and searches for biases in an agency’s AI all depend to
32 a large extent on transparency about the AI technique and how agency personnel interact with it.
33 Although this Statement addresses considerations separately for purposes of exposition, agencies
34 should be mindful about the extent to which addressing one consideration might affect others.

35 With these overarching principles in mind, the Administrative Conference believes that
36 the following, non-exhaustive set of considerations can provide a useful starting point for
37 agencies interested in making greater or different use of AI tools.

38 *Transparency*

39 It is important that agencies take transparency concerns seriously in developing and
40 deploying AI. The appropriate level or nature of transparency in an agency’s AI systems will

Commented [A3]: A commenter raises the following concern:
“A focus on ‘transparency’ must emphasize that, because of technical, legal, and policy reasons, transparency may be necessarily limited.... The general takeaway should be that each agency will approach the question of transparency differently based on the specific AI system or use at issue, and that transparency is not absolute.”

² In performing government functions, agencies will likely use both standalone AI techniques as well as AI techniques embedded in commercial products. Agencies may have significantly less control over the design, development, and, to some extent, use of embedded AI technologies, and thus different considerations may apply to such technologies.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

41 inevitably depend on context, including any applicable laws and policies governing disclosure of
42 information. For example, when an agency deploys AI in adjudication, the need to give regulated
43 parties a full explanation of the decisionmaking process may require a high degree of
44 transparency from the agency regarding how the AI functions. By contrast, when an agency uses
45 AI to make or assist in making enforcement decisions, the agency’s legitimate interest in
46 preventing gaming or adversarial learning by regulated parties could militate against the
47 agency’s publicly providing too much information (or specific types of information) about the
48 AI’s processes. Agencies should be sensitive to such context-based distinctions in evaluating
49 transparency in an AI.

50 Among other things, agencies should carefully consider to whom they should be
51 transparent and for what purposes. For instance, depending on the nature of its operations, an
52 agency might prioritize transparency to the public, courts, or its own officials. And the agency
53 might prioritize transparency in the service of various goals, such as legitimizing its AI tools,
54 facilitating internal and external review of its AI-based decisionmaking, or ensuring the smooth
55 coordination of its activities. Different types of AI are likely to serve some of these concerns
56 better than others, something agencies should bear in mind when thinking about whether and
57 how to employ particular AI techniques.

58 Transparency in AI can enhance accountability by yielding more reliable, verifiable, and
59 trustworthy agency decisionmaking processes. To that end, in choosing and using AI techniques,
60 agencies should be cognizant of the degree to which a particular AI technique is interpretable
61 and understandable by agency experts and explainable to appropriate stakeholders, including the
62 general public. At the same time, however, agencies should be aware that there may exist
63 tradeoffs between explainability and accuracy in AI. The appropriate balance between
64 explainability and accuracy will depend on the agency’s unique circumstances, including its
65 priorities when it comes to transparency.

66 It is also important for agencies to think about how they will explain decisions made by
67 their AI techniques. Technologies for explaining AI decisionmaking processes—commonly
68 called “explanatory AI” or “xAI”—are rapidly evolving. Different types of xAI offer different



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

69 ways of explaining certain AI techniques and some of those ways of explaining AI techniques
70 may be more useful than others, depending on the technique involved and the level of
71 explanation required. In deciding what type of AI techniques to employ, therefore, agencies
72 should carefully consider the extent to which different techniques can be explained, and be
73 mindful of the fact that the tradeoff between explainability and accuracy may often militate in
74 favor of choosing simpler AI models.

75 In thinking about transparency issues in AI, agencies should also consider questions
76 about intellectual property when procuring or using AI. When an agency's AI technique relies on
77 proprietary technologies or algorithms that the agency does not own, intellectual property
78 protections can limit the agency's and the public's access to information that may allow it to
79 understand or explain the AI technique. For that reason, agencies should work with outside
80 providers to ensure they can use a given AI technique in a transparent and explainable fashion.
81 Because intellectual property limitations can adversely affect an agency's ability to be
82 transparent about the AI techniques it uses, agencies may wish to consider adopting a preference
83 for non-proprietary technologies, except when there is a clear and well-documented case that the
84 advantages of using such technologies outweighs the costs in lower transparency and public
85 credibility that may accompany reliance on proprietary technologies.

Judicial Review of AI Applications

87 Those who are affected by agency uses of AI may file lawsuits challenging the agency's
88 actions. For instance, when an agency adopts an AI technique that narrows the discretion of
89 agency personnel or fixes or alters the legal rights and obligations of people subject to the
90 agency's action, affected people or entities might allege that the AI technique is a legislative rule
91 and might sue the agency if the agency neglected to put the rule through the notice and comment
92 processes the Administrative Procedure Act (APA) ordinarily requires for legislative rules.

93 Determining whether an AI technique operates as a legislative rule—that is, whether it has the
94 “force of law” by virtue of narrowing agency officials’ discretion or affecting the legal rights and
95 obligations of regulated parties—may require analysis of the technical system and the interface
96 design and policies guiding the system’s use by agency personnel. If such an analysis shows that

Commented [A4]: A commenter makes the following suggestion:

“Principles on the Review of AI Decisions should not be limited only to judicial review. While the adjudication/litigation context is important, in most circumstances, reviews of AI systems/use will be conducted by internal and external government entities.

For instance, OIG, GAO, and congressional oversight and reviews are, in many respects, more prevalent and more detailed. While briefly touched on in the AI Audit section below, the reviews conducted by these entities typically go beyond traditional concepts of auditing, and operate more as an oversight role that impacts agency operations, sometimes more so than litigation.”



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

97 the AI technique might operate with the force of law, agencies should treat choices surrounding
98 the AI technique as they would the formulation of a substantive rule, by prioritizing the need for
99 technical expertise, public input, and reasoned, thoroughly-documented decisionmaking.

100 If, on the other hand, the agency’s analysis shows that the AI technique is merely being
101 used to support, but not constrain, agency actions, the agency might elect not to treat the AI
102 technique as a legislative rule, and might instead treat it as a manifestation of internal agency
103 procedure and management. Before making that determination, agencies should ask themselves
104 whether the AI technique operates within the zone of discretion left by the rule, and whether it
105 leaves the substance of the rule intact. To the extent an agency can answer those questions in the
106 affirmative, it is concomitantly less likely that litigation demanding compliance with notice and
107 comment requirements will succeed. Even when an AI technique does not operate as a legislative
108 rule, however, prioritizing the need for technical expertise, public input, and reasoned,
109 thoroughly-documented decisionmaking will often be advisable.

110 Agency use of AI techniques might also lead to litigation in cases where parties affected
111 by an agency’s use of an AI technique allege that the AI technique violated their rights to
112 procedural due process. Courts would analyze such challenges under the three-part balancing
113 framework set forth in *Mathews v. Eldridge*, which inquires after (1) “the private interest that
114 will be affected by the official action;” (2) “the risk of an erroneous deprivation of such interest
115 through the procedures used, and the probable value, if any, of additional procedural
116 safeguards;” and (3) “the Government’s interest, including the function involved and the fiscal
117 and administrative burdens that the additional or substitute procedural requirement would
118 entail.”³ Some AI techniques will fare well under that framework, particularly when they deliver
119 more accurate decisionmaking at lower costs to agencies. Even if an AI technique does not
120 violate people’s rights to due process as a matter of law, however, the concerns that animate due
121 process—like reasoned decisionmaking, transparency, and predictability—will often be highly
122 relevant, as a policy matter, to agencies’ decisions about employing different AI techniques.

Commented [A5]: Some commenters have supported language along these lines, others have questioned whether the committee has a basis for this.

³ 424 U.S. 319, 335 (1976).



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

123 Agency uses of AI may also result in lawsuits alleging violations of the APA’s
124 prohibition on agency actions that are “arbitrary, capricious, an abuse of discretion, or otherwise
125 not in accordance with law.”⁴ In adjudicating such lawsuits, courts will consider whether the AI
126 decision “relied on factors on which Congress has not intended [the agency] to consider, entirely
127 failed to consider an important aspect of the problem, offered an explanation for its decision that
128 runs counter to the evidence before the agency, or is so implausible that it could not be ascribed
129 to a difference in view or the product of agency expertise”⁵ Complying with that standard will
130 presumably require that agencies be able to (1) articulate what the relevant factors are, (2)
131 demonstrate that the AI technique is designed to take them into account, and (3) that the AI
132 technique accounts for the relevant factors in a way that is reasonable under the circumstances.⁶

133 *Performance Evaluation and Oversight*

134 Because AI is developed by fallible humans, flaws in AI techniques are inevitable,
135 especially as the techniques become more complex. In evaluating and overseeing their AI
136 techniques, however, agencies should not necessarily let the discovery of defects dissuade them
137 from continuing to use the techniques. Agencies should be mindful that the relevant evaluation is
138 always a comparative one; the ultimate question is whether the AI technique fares better than
139 human decisionmakers.

140 To identify and correct for flaws in their AI techniques, agencies should rely on a variety
141 of evaluation and oversight mechanisms, both internal and external. Among these, agencies
142 should ordinarily prioritize methods that address the types of systematic errors that sometimes
143 plague AI techniques.

⁴ See 5 U.S.C. § 706(2)(A).

⁵ *Motor Vehicles Mfrs. Ass’n v. State Farm Mut. Auto. Ins. Co.*, 463 U.S. 29, 43 (1983). This “hard look” standard contrasts with the standard applied in cases like *Hayward v. Department of Labor*, 536 F.3d 376, 377 (5th Cir. 2008), whereby courts decline to consider whether an AI considered relevant factors.

⁶ The extent to which courts do or ought to probe agencies’ rationales in relying on something like an AI technique is the subject of considerable debate. This Statement does not purport to address that debate, save for acknowledging that it exists.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

144 Among external evaluation and oversight mechanisms, litigation may not be especially
145 effective at correcting the sorts of systematic errors that are common in different AI techniques.
146 To be sure, some litigants may challenge the use of an AI technique across a whole class of
147 cases, and such litigation might serve as an important external check on systematic flaws in AI.
148 But many litigants will only seek relief in the form of a remedy in their individual cases, so that
149 those lawsuits may not be an especially effective check on systematic errors. Moreover,
150 traditional administrative law doctrines can make it difficult for litigants to obtain review many
151 types of agency decisionmaking that might be plagued by systematic errors. Deferential
152 standards of review limit the extent to which courts can require agencies to use particular
153 methods for making decisions. And doctrines about agencies' enforcement discretion effectively
154 hive off most enforcement decisions from judicial review. Consequently, when it comes to
155 correcting systematic errors resulting from AI techniques, agencies should look beyond judicial
156 review in individual cases.

157 That may be particularly so given that many considerations relating to agencies' uses of
158 AI that are not cognizable on judicial review may nevertheless implicate important
159 administrative law considerations, such as reasoned decisionmaking, fairness, and providing an
160 opportunity to be heard. By employing rigorous and thoughtful evaluation and oversight of the
161 selection and application of chosen AI techniques, agencies can build public support for
162 themselves and their uses of AI techniques and increase perceptions of such techniques as
163 legitimate and fair.

164 Regular evaluation and oversight of AI techniques, throughout the techniques' lifespans,
165 is especially important. Agencies should account for the need for such evaluation and oversight
166 early on, even during the procurement process, so that they can negotiate agreements that allow
167 them to avoid the intellectual property issues that might make it difficult for agencies and others
168 to evaluate and oversee AI techniques. Once an agency has procured an AI system, it should
169 establish a plan for regularly evaluating the AI system, particularly if the system or the
170 circumstances in which it is deployed are liable to change over time. For instance, if an agency's
171 AI technique uses unlocked models—that is, models that continuously update themselves—the



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

172 consequence will be a dynamic algorithm, constantly changing over time. With AI techniques
173 based on unlocked models, review and explanation of the algorithm’s functioning at one stage of
174 development or use may quickly become outdated due to substantive changes in the algorithm’s
175 underlying models. These changes can operate in two directions: they can cause problems or
176 eliminate them. For the agency to know how its AI technique is functioning, however, regular
177 evaluation and oversight is imperative so that the agency can become aware of deficiencies in the
178 technique before those deficiencies create widespread problems.

179 Evaluation and oversight of agencies’ uses of AI will often be internal, but it is important
180 that agencies also consider how they can effectively solicit external input about their AI. In
181 addition to providing agencies with a more diverse set of perspectives for identifying potential
182 problems with the AI, external evaluations and oversight of agency AI uses can also help
183 legitimize those uses in the eyes of both those who are affected and those who will conduct
184 oversight reviews of the agency’s work. Agencies might obtain such evaluations by contracting
185 for them. Alternatively, agencies could submit their AI techniques to public comment and
186 thereby obtain outside input on them.

187 In evaluating their uses of different kinds of AI, agencies should beware that AI systems
188 can affect how agency staff do their jobs, particularly as staff grow to trust and rely on the
189 systems. Automation bias, the tendency of non-experts to be overly deferential to decisions made
190 by computers, is a well-documented problem, particularly in circumstances where agency
191 decisionmakers face heavy workloads. In addition to evaluating and overseeing their AI
192 technologies, therefore, agencies should pay attention to how agency staff uses those
193 technologies and whether the technologies might be impairing staff’s performance of other job
194 functions. Along the same lines, agencies should consider how AI might affect their staffing
195 needs.

196 Agencies should ensure that they have plans for evaluating their AI techniques and that
197 officials at the agency respond to what is learned from the evaluation. The Administrative
198 Conference recognizes that the resource implications for implementing life-cycle AI evaluation
199 are potentially substantial. However, given the increasing importance of AI to a wide range of

Commented [A6]: A commenter suggests that what’s missing before this paragraph is guidance about internal quality assurance systems distinct from procured AI — ideally something similar to the range of quality assurance / quality improvement programs in the human context. The commenter cites a piece co-authored by one of the project consultants. *See* <https://www.stanfordlawreview.org/print/article/due-process-and-mass-adjudication/>

What’s important, the commenter points out, is that we have an internal administrative and internal due process to make sure that AI systems are developed and maintained and benchmarked. The statement might make that point by including a paragraph about internal oversight that integrates the next paragraph (about the sociotechnical dimension in this) before external oversight.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

200 agencies, such investments are critical to enabling AI to meet challenges now and in the future.
201 Agencies should take care to allow sufficient time, resources, and structure to ensure that reviews
202 of AI uses are adequate.

203 *Unintended Bias*

204 As with human decisionmaking, AI techniques can be biased. Although AI can help
205 agencies identify and reduce the impact of human biases, it can also unintentionally create or
206 exacerbate those biases by encoding and deploying them at scale. In deciding whether and how
207 to deploy an AI technique, therefore, agencies should carefully evaluate whether the technique is
208 likely to create or exacerbate unintended biases more than the human decisionmaker it might
209 replace. In conducting that evaluation, agencies should also be cognizant that biases have
210 different consequences in different contexts, meaning the consequences of certain types of AI-
211 related biases will be different depending on the nature of the decisions the AI is tasked with
212 making.

213 Unintended biases can creep into an AI tool or system in several ways, including:

- 214 • Using training data that includes biased human decisions or reflects historical or
215 social inequities;
- 216 • Defining an AI technique’s target variable and the associated class labels in ways
217 that reflect biases;
- 218 • Developing models that reflect biases, such as in the choice of attributes an AI
219 technique observes and folds into its analysis;
- 220 • Using flawed data sampling, in which groups are over- or under-represented in
221 training data; and
- 222 • Applying an AI technique in one setting even though it was trained on data from a
223 materially different setting.

224 Biases from these and other sources can propagate over time through feedback loops,
225 whereby the use of a biased AI technique creates more biases, which are then reflected in the
226 data the technique uses to make future decisions. In seeking out biases in AI techniques, agencies



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

227 should be mindful of the interdependence of AI and the models, metrics, and data that underpin
228 it.

229 There may be some cases where identifying biases in given AI techniques is especially
230 difficult, for instance when protected attribute information is not directly available. In such cases,
231 rather than abandoning attempts to identify and mitigate biases, agencies should try to make use
232 of other available methods, such as record linkage or imputation. In deciding which alternative
233 methods to attempt, agencies should be mindful of their obligations to explain their
234 decisionmaking regarding AI—to courts, Congress, and the public.

235 It is possible that certain interventions that correct for systematic biases in AI techniques
236 will lead to challenges on equal protection grounds. Despite those risks, agencies should work
237 diligently to identify and mitigate unintended biases in the AI techniques they use. Doing so will
238 require agencies and their personnel to stay up to date on developments in the field of AI,
239 particularly on algorithmic fairness; establish processes to ensure that people with diverse
240 perspectives are able to inspect the AI and its decisions for hints of unintended bias; test
241 algorithms in regimes resembling the ones in which they will be used; and make use of technical
242 and external tools, such as red teams and third-party audits, to supplement internal agency
243 processes for evaluating bias risks in AI systems.

244 *Technical Capacity and Contracting*

245 Implementing AI techniques may help agencies conserve scarce resources. But it can also
246 require a major investment of human and financial capital—both in procuring or developing an
247 AI system and also in maintaining it. Agencies should carefully evaluate the short- and long-term
248 costs and benefits of an AI before committing to it.

249 The costs of developing an AI tool will depend in part on whether the agency creates it
250 internally or procures it from an external source. There are different benefits and drawbacks to
251 each approach. The decision to buy an AI system from an external source might allow the
252 agency to acquire a more sophisticated tool than it could design on its own, to acquire a tool in a
253 more efficient timeframe, and to save some of the up-front costs associated with developing the



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

254 technical capacity needed to design an AI system. The decision to create an AI tool within the
255 agency, by contrast, may yield a tool that is better tailored to the agency’s particular tasks and
256 more compliant with the agency’s policy aims. Creating an AI system within the agency can also
257 provide an impetus for the agency to develop internal technical capacity, which can yield
258 benefits during the lifetime of the AI system and can also be exploited in other technological
259 contexts.

260 The appropriate answer to the “make or buy” question regarding AI will ultimately
261 depend on the agency’s unique circumstances. To ensure an informed decision, an agency should
262 involve personnel from all of the relevant offices—including at least personnel involved in the
263 agency’s legal, policy, procurement, financial, human resources, and technology offices—are
264 consulted, because decisions about implementing or modifying a particular AI technique can
265 have implications for all of those agency functions. The same is true with respect to decisions
266 about how to structure agreements for procuring or modifying AI techniques. Agencies can
267 minimize the risk of unforeseen problems by making sure that all of their affected units are heard
268 in the course of making decisions about the agency’s uses of AI.

269 Because it takes a high level of technical sophistication to make an intelligent decision
270 about whether to make or buy AI, agencies facing that choice should ensure they have access to
271 the relevant technical expertise. Given the relative dearth of experts in the emerging field of AI,
272 as well as the field’s ongoing and rapid development, agencies should be prepared to expend the
273 financial and human resources to make sure their decision about obtaining AI is well-informed
274 by people with the requisite knowledge.

275 In some cases, agency personnel will, themselves, lack the skill to build an AI system that
276 meets the agency’s needs. In those cases, agencies should consider other means of quickly
277 expanding their technical expertise, including by relying on tools like the Intergovernmental
278 Personnel Act, prize awards, or cooperative research and development agreements with private

Commented [A7]: Some commenters have suggested that this language is likely to become dated.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

279 institutions or universities.⁷ Additionally, agencies should prioritize cultivating an AI-ready
280 workforce, including through recruitment and training efforts that emphasize AI skills.

281 An agency that elects to purchase an AI tool from an external source rather than building
282 it in-house should still invest in developing their internal capacity to make sound decisions about
283 the type of AI the agency requires, how to integrate that AI system into the agency's information
284 technology resources, and how to maintain the system once operational. It is important, for
285 example, that the agency personnel who negotiate the purchase agreement know enough about
286 AI so that they can reasonably foresee what information the agency might need about the AI's
287 functioning in the future. Armed with such knowledge, agency personnel should make sure that
288 any agreement they sign to procure AI allows them access to the information they will need to
289 provide about the AI.

290 Agencies that are uncertain about whether they have the internal resources to make fully
291 informed decisions about building or contracting for AI should avail themselves of other
292 government offices that exist to help with such decisions.⁸

Delegation and Accountability

294 Agencies that use AI may also confront issues involving purported delegations of agency
295 authority. An agency that relies on an AI technique might, for instance, be thought of as
296 delegating its own authority to the technique or to the designer of the technique. From a legal
297 standpoint, such purported delegations seem unlikely to create legal difficulties. They are akin to
298 an agency's using a model or formula as a basis for decisions. So long as the agency can explain

⁷ For example, the Food and Drug Administration's Centers of Excellence in Regulatory Science and Innovation (CERSIs) are collaborations between FDA and academic institutions to advance regulatory science through innovative research, training, and scientific exchanges. They have yielded considerable AI-related benefits in the medical device field.

⁸ Within the General Services Administration, for example, the office called 18F routinely partners with government agencies to help them build and buy technologies. Similarly, the United States Digital Service has a staff of technologists whose job is to help agencies build better technological tools. While the two entities have different approaches—18F acts more like an information intermediary and the Digital Service serves as an alternative source for information technology contracts—both models could be effective, if built out, in aiding agencies with obtaining, developing, and using different AI techniques.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

299 how a model works and why the agency reasonably chose to use that model, such agency
300 practices have never been thought of as implicating nettlesome delegation questions under the
301 law. The same holds true for agency uses of AI techniques.

302 The more interesting delegation questions confronting agencies that use AI are ones the
303 law does not yet resolve. In particular, agency uses of AI techniques can raise questions about
304 the extent to which regulated parties are entitled to human decisionmaking, rather than
305 computerized decisionmaking. While courts have yet to interpret the APA or other statutes to
306 forbid computerized decisionmaking, values of transparency, accountability, and due process
307 will, in some cases, weigh in favor of at least providing human oversight to check and confirm
308 AI determinations.

309 A related type of delegation question involves internal and informal administrative law.
310 When an agency creates or uses an AI technique, the effect of doing so can be to shifts discretion
311 around within the agency. In some cases, agency use of AI techniques can increase managerial
312 control by systematizing and regularizing certain determinations. At the same time, agency use
313 of AI techniques can also have the effect of increasing the power of agency technologists,
314 particularly as the AI in question becomes more complex and less intelligible to non-experts.
315 Agencies should be aware of these potential shifts of authority and, where appropriate, take steps
316 to ensure that appropriate officials have the knowledge and power to be accountable for
317 decisions that are made or aided by AI techniques.

318 *Data Collection and Standardization*

319 To enable their AI, agencies must collect or obtain data, often in vast quantities. In
320 collecting and organizing that data, agencies should consider (1) whether they have the
321 appropriate data set, meaning it reflects conditions similar to the ones the AI will address in
322 practice; (2) whether the data is in a form that is usable without an excessive amount of labor,
323 and (3) how the agency ensures that the data it collects is maintained and linked to the AI.

324 Agencies stand to benefit greatly from large and broad pools of data collected and stored
325 in usable forms. To enhance the availability of such data, agencies may wish to consider

Commented [A8]: Some commenters have suggested adding a footnote to this section (perhaps here) referring to something like the IC's framework for AI, simply as a point of comparison. See <https://www.dni.gov/index.php/features/2763-principles-of-artificial-intelligence-ethics-for-the-intelligence-community>.



ADMINISTRATIVE CONFERENCE OF THE UNITED STATES

326 collaborating to advance the **Federal Data Strategy**, which encourages agencies to “[a]dopt or
327 adapt, create as needed, and implement data standards within relevant communities of interest to
328 maximize data quality and facilitate use, access, sharing, and interoperability.” Given agencies’
329 unique data needs, such collaboration in standardizing data will not always be practicable, but it
330 is something to which agencies might aspire.

331 *Privacy*

332 Agencies have a responsibility to protect privacy when an AI system creates, collects,
333 uses, processes, stores, maintains, disseminates, discloses, or disposes of personally identifiable
334 information. Although this responsibility is not unique to AI, it adds an additional layer of
335 analysis, which may impact the issues described above. In a narrow sense, it means complying
336 with specific legal and policy requirements related to, for example, transparency, due process,
337 accountability, and information quality and integrity, established by the Privacy Act of 1974,
338 Section 208 of the E-Government Act of 2002, and other laws and policies.⁹ In a broader sense,
339 it means recognizing and appropriately managing any privacy risks present in an AI system,
340 especially those of a unique or heightened nature. Agencies should consider privacy risks
341 throughout the entire system development life cycle and assess those risks, as well as associated
342 controls, on an on-going basis. Office of Management and Budget (OMB) Circular A-130,
343 *Managing Information as a Strategic Resource* (July 28, 2016), National Institute of Standards
344 and Technology (NIST) Special Publication SP-800-37 revision 2, *Risk Management Framework*
345 *for Information Systems and Organizations: A System Lifecycle Approach for Security and*
346 *Privacy* (Dec. 2018), and related documents establish a Risk Management Framework for
347 agencies to utilize when implementing an AI system.

348 *Potential Projects for the Administrative Conference Regarding AI*

Commented [A9]: A commenter suggests this might be a good place to drop a footnote on proposals for a National Research Cloud that enables such a data strategy for AI. See <https://www.nextgov.com/emerging-tech/2020/06/congress-seeks-creation-national-research-cloud-artificial-intelligence/165954/>

⁹ See, e.g. 5 U.S.C. § 552a(e), (g), & (p); 44 U.S.C. § 3501 note.